

Docket No: POU920010169US1

METHOD AND APPARATUS FOR
OBTAINING MULTIPLE PORT
ADDRESSES BY A FIBRE CHANNEL
FROM A NETWORK FABRIC

APPLICATION FOR
UNITED STATES LETTERS PATENT

Express Mail Label No: ET089525673US

Date of Deposit: December 3, 2001

I hereby certify that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee Service under 37 CFR 1.10 on the date indicated above and is addressed to Box Patent Application, U.S. Patent and Trademark Office, P.O. Box 2327, Arlington, VA 22202.

Susan L. Nelson

I N T E R N A T I O N A L B U S I N E S S M A C H I N E S
C O R P O R A T I O N

METHOD AND APPARATUS FOR OBTAINING MULTIPLE PORT ADDRESSES BY A FIBRE CHANNEL FROM A NETWORK FABRIC

BACKGROUND OF THE INVENTION:

[0001] The present invention is related to establishing a data path from a channel to an I/O adapter in a data processing system having a network fabric, and is more particularly related to obtaining multiple port addresses by a fibre channel directly from the network fabric.

[0002] A host Fibre Channel adapter supporting multiple logical partitions (LPARs) or other entities which are executing the Fibre Channel/SCSI (FCP) protocol must have a unique Fibre Channel N_Port address identifier for each LPAR supported. However, there is no existing method by which an adapter with a single Fibre Channel N_Port can obtain multiple N_Port address identifiers directly from the fabric F_Port.

[0003] The need to obtain multiple (virtual) N_Port identifiers is a new requirement which arises when multiple host LPARs sharing the same Fibre Channel adapter, execute the SCSI/Fibre Channel protocol (FCP) through a single N_Port. It is advantageous to minimize the number of adapters to reduce cost and to maximize adapter utilization. Others have not solved this problem previously.

[0004] There are other protocols which result in an N_Port having more than one ID, such as when an N_Port receives multicast frames (sent to a multicast address) as well as frames sent to its own N_Port address. These other protocols are specifically designed for unique purposes (such as multicast), however, and they require the presence of supporting "servers" in

the fabric (e.g. the Multicast Server). They also have limitations which result from the specific function provided. For example, multicast addresses are shared by all N_Ports in the multicast group, but an address used by a logical partition must only be useable by that logical partition.

[0005] Another potential method of obtaining another N_Port ID could be achieved by allowing the N_Port to log in multiple times with the fabric, but this would require initialization of flow-control buffers, and would therefore disrupt the operation of preexisting virtual N_Ports with other N_Port Ids.

[0006] U.S. Patent No. 5,276,813 issued January 4, 1994 to Elliott et al. for ACQUIRING ADDRESSES IN AN INPUT/OUTPUT SYSTEM, discloses a computer Input/Output system in which link-level facilities issue an acquire link address frame when initially coming on line. A dynamic switch receiving the frame then assigns a link address to the link-level facility.

[0007] U.S. Patent No. 5,420,988 issued May 30, 1995 to Elliott for ESTABLISHING LOGICAL PATHS THROUGH A SWITCH BETWEEN CHANNEL AND CONTROL UNITS IN A COMPUTER I/O SYSTEM, discloses a mechanism for assigning multiple logical path identifications with a single physical path.

[0008] U.S. Patent No. 6,084,859 issued July 4, 2000 to Ratcliff et al. for INTERNET PROTOCOL ASSISTS USING MULTI-PATH CHANNEL PROTOCOL, discloses an apparatus for allowing any initiating host to establish communications with any receiving host in a computing network using a multi-path channel communication protocol.

SUMMARY OF THE INVENTION:

[0009] In order to obtain multiple N_Port identifiers, the N_Port first logs in with the fabric by sending a "Fabric Login" (FLOGI) extended link service (ELS) command to the attached F_Port using a source address of all zeros. This step is a normal initialization procedure, as is well known in the art, and is performed by almost all implementations. Upon completion of this step, the N_Port has been assigned its first N_Port address identifier, and service parameters have been transferred. After fabric login in complete, the fabric prepares itself to assign additional N_Port identifiers, and "implicitly" logs in these additional N_Port identifiers. The additional identifiers will be assigned upon the receipt of FDISC as described below.

[0010] In order to obtain another N_Port address identifier, the N_Port sends an FDISC ELS command using either a source address identifier of zero or, if known, the new source address identifier using identical service parameters as provided in the original FLOGI. The FDISC ELS is used instead of additional FLOGIs to avoid disruption of the operating environment.

[0011] When the N_Port sends the FDISC ELS to the fabric the FDISC provides following functions:

1. It informs the Fabric that a new (virtual) N_Port (and logical partition) exists behind the physical port.
2. It provides the means for the virtual N_Port to transfer a unique Port_Name to the fabric.
3. It provides a signal to the fabric to validate and assign the virtual N_Ports' new N_Port ID, and allows both the fabric and the virtual N_Port to begin normal frame reception and transmission.
4. It provides a signal which causes the fabric to update the name server if necessary or any other database maintained within the fabric.

[0012] The use of the FDISC ELS to obtain a new N_Port ID has no effect with ongoing operations of preexisting virtual N_Ports, does not violate existing standards, and does not require the presence of specialized servers within the fabric.

BRIEF DESCRIPTION OF THE DRAWINGS:

[0013] These and other objects will be apparent to one skilled in the art from the following detailed description of the invention taken in conjunction with the accompanying drawings in which:

Fig. 1 is a schematic diagram of a data processing network having a server which has at least one N_Port and a network fabric;

Fig. 2 is a flowchart of the procedure followed by the N_Port of the system of Fig. 1 for requesting an N_Port address form the fabric; and

Fig. 3 is a flowchart of the procedure followed by the fabric of Fig. 1 for providing addresses to the N_Port.

DESCRIPTION OF THE PREFERRED EMBODIMENT:

[0014] Fig. 1 is a schematic diagram of a data processing network usable with the present invention. The data processing network includes a server or host 100 which may be, for instance, an IBM zSeries 900 server. The server 100 includes an operating system which provides for supporting multiple logical partitions (LPARs) 102a-102n. Each LPAR 102 communicates with I/O devices through channel adapters, as is well known. The server 100 is

connected to a fabric 110 by a channel adapter 104, with the fabric 110 being connected to one or more controllers 120a-120m controlling I/O devices (not shown) as is well known. The connection and sending of data between the server 100, the fabric 110, and the controllers 120 may, for instance, be as described in proposed standard Fibre Channel Framing and Signaling (FC-FS Rev. 1.40) NCITS Project 1331-D, October 15, 2001.

[0015] The adapter 104 is, in the exemplary embodiment, hardware having a software interface for communicating with the LPARs 102, and includes an N_Port 106 for connecting to an F_Port 112 of the fabric 110. Although a single port, the N_Port 106 recognizes multiple addresses, one address for each of the partitions 1-n (LPARs 102a-102n). Once the address for partition 102 is established, communication by, for instance, a controller to that particular partition 102 may be effected by communicating with the partition's assigned address. A name server 114 is provided in the fabric 110 which includes an address table 116. This address table includes, among other things, an identification of each partition 102, and its corresponding N_Port address. Thus, in order to communicate with a partition, the name server 114 may be accessed to determine the N_Port address corresponding to that partition 102, and data is then sent to that N_Port address. It will be understood that each F_Port 112 in the fabric 110 acts to receive and transmit data and commands between the fabric 110 and the channel adapter 104 or the controllers 120a-120m, as is well known.

[0016] Fig. 2 is a flow diagram of the procedure followed by the N_Port 106 of a fibre channel to obtain an N_Port address from the fabric 110. The procedure starts at 200. If at 201, this is the first address for the N_Port 106, a Fabric Login extended link service (FLOGI ELS) command is issued to the fabric 110. The FLOGI ELS command includes an identification for the

partition 102 requesting the address, and has a source address of all zeros to indicate to the fabric 110 that an N_Port address identifier is being requested. At 203, the address assigned by the Fabric 110 is received. The FLOGI command is sent in this embodiment on behalf of the first partition 102 requesting an address. This first request may always be assigned by a selected one of the partitions, or may be the first partition needing an address, or may be selected by, for instance, a round robin scheme, as may be desired.

[0017] If this is not the first address being selected for a partition 102, at 204, a Fabric Discovery Extended Link Service (FDISC ELS) command is issued using either the source address identifier of zero, or, if known, the new source address identifier using identical service parameters as provided in the original FLOGI command. At 205, the N_Port 106 receives the address assigned or confirmed by the fabric 110 for use with the requesting partition 102.

[0018] Fig. 3 is a flow diagram of the procedure followed by the fabric 110 in assigning the address identifications requested in the flowchart of Fig. 2, and starts at 200. At 301, the fabric 110 receives the ELS command with a partition ID. The partition ID may be the worldwide partition number, or any other identification scheme to identify the partition to be associated with the N_Port address identification. At 302, it is determined if the command is a FLOGI command. If yes, at 303 an address table 116 is established in the name server 114 for the N_Port 106. At 304, the first address identification is assigned for this N_Port 106. At 305, the address identification is recorded in the table 116, along with the partition identification, and other parameters needed for the communications protocol to be used to transfer commands and data between the N_Port 106 and the

controllers 120a-m. At 306, the address is returned to the N_Port 106.

[0019] If the ELS is an FDISC command at 307, the next address is assigned or confirmed at 308. If a new source address identification is supplied by the FDISC ELS command, that address identification, if acceptable, is used. If an address is not supplied, the next available address is assigned by 110 in accordance with a desired scheme which insures that duplicate numbers are not assigned. In addition at 308, if the partition ID already has an address identification in the table 116, the address identification is updated by the identification in the FDISC ELS command. Thus, the FDISC ELS command may be used to request an address be assigned, may have a proposed address identification confirmed, or may update an old address identification with a new address identification. Then at 305, the partition identification, address identification, and other parameters are recorded in the address table 116, and at 306, the assigned, confirmed, or updated address identification is returned to the N_Port 106. It will now be understood that normal frame reception and transmission can begin. The controllers will see "n" different adapters, one adapter for each partition, but there will be only one channel adapter N_Port 106 with "n" different addresses.

[0020] While the preferred embodiment of the invention has been illustrated and described herein, it is to be understood that the invention is not limited to the precise construction herein disclosed, and the right is reserved to all changes and modifications coming within the scope of the invention as defined in the appended claims.